

FILTERED ROUTER ALERT

HOP-BY-HOP OPTION

INVENTOR:

Glenn Morrow

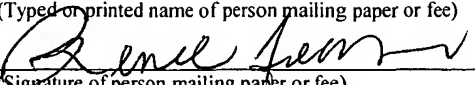
"Express Mail" mailing label No. EU263079705US

Date of Deposit: July 14, 2003

I hereby certify that this paper or fee is being deposited with the United States Postal Service, "Express Mail" service under 37 C.F.R. 1.10, on the date indicated above and is addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450

Renee Fears

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)

FILTERED ROUTER ALERT HOP-BY-HOP OPTION

5

Related Application Data

10 This application is related to Provisional Patent Application Serial
No. 60/398,206 filed on July 24, 2002, and priority is claimed for this earlier filing under 35 U.S.C. § 120. The Provisional Patent Application is also incorporated by reference into this utility patent application.

15

Technical Field of the Invention

A message protocol for identifying an information packet for increased inspection and handling.

BACKGROUND OF THE INVENTION

The Internet, like so many other high tech developments, grew from research originally performed by the United States Department of Defense. In the 1960s, the military had accumulated a large collection of incompatible computer networks. Computers on these different networks could not communicate with other computers across their network boundaries.

In the 1960s, the Defense Department wanted to develop a communication system that would permit communication between these different computer networks. Recognizing that a single, centralized communication system would be vulnerable to attacks or sabotage, the Defense Department required that the communication system be decentralized with no critical services concentrated in vulnerable failure points. In order to achieve this goal, the Defense Department established a decentralized standard communication protocol for communication between their computer networks.

A few years later, the National Science Foundation (NSF) wanted to facilitate communication between incompatible network computers at various research institutions across the country. The NSF adopted the De-

fense Department's protocol for communication, and this combination of research computer networks would eventually evolve into the Internet.

Internet Protocol and Packet-Based Communication

5 The Defense Department's communication protocol governing data transmission between different networks was called the Internet Protocol (IP) standard. The IP standard uses discrete information packets, sometimes called datagrams, to communicate between different computers and other devices and networks over the Internet. The IP standard has been widely adopted for the transmission of discrete information packets
10 across network boundaries. In fact, most telecommunication networks operate using information packets to transmit data to linked communication devices. The IP standard or similar packet-based communication protocols govern communications on these networks as well as the Internet.

Packet-based communication protocols depend on destination and
15 source address data found in an address header for routing over a communication network. This type of network is sometimes referred to as a packet-switched network, because each information packet's path through the network is controlled by switching or routing decisions based on the address data found in the packet's address header. For these types of
20 communication networks, the communication protocols operate to estab-

lish an end-to-end connection using individual information packets or datagrams each following individually set routes. Since each information packet is individually routed over the network without a fixed path or route, these networks are also characterized as connectionless networks.

5 In a typical information packet-based communication scenario, data is transmitted from an originating communication device on a first network across a transmission medium to a destination communication device on a second network. During transmission, transit routers on the network process the information packet address header to route the indi-
10 vidual information packets. After receipt at the destination device, the destination communication device decodes the transmitted information into the original information transmitted by the originating device according to the applicable communication protocol.

Addressing and Routing

15 A communication device operating on an information packet-based network is assigned a unique physical address. For IP-based networks, this address is referred to as an IP address. The IP address can include:
(1) a network ID and number identifying a network, (2) a sub-network ID number identifying a substructure on the network, and (3) a host ID num-
20 ber identifying a particular computer on the sub-network. A header data

field in the information packet will include source and destination addresses. The IP addressing scheme imposes a consistent addressing scheme that reflects the internal organization of the network or sub-network. Other addressing protocols use address headers and similar addressing mechanisms to route information packets.

A router is used to regulate the transmission of information packets into and out of the communication network. Routers interpret the logical address contained in information packet headers and direct the information packets to the intended destination. Information packets addressed between communication devices on the same network do not pass through a router on the boundary of the network, and as such, these information packets will not clutter the transmission lines outside the network. If data is addressed to a communication device outside the network, the router on the network boundary forwards the data onto the greater network.

Network communication protocols define how routers determine the transmission path through a network and across network boundaries. Routing decisions are based upon information in the address header and corresponding entries in a routing table maintained on the router. A routing table contains the information for a router to determine whether to ac-

cept an information packet on behalf of a device or pass the information packet onto another router.

Routing that involves processing at a router and then forwarding in a hop-by-hop manner from one router to the next is referred to as hop-by-hop routing. At each point in the routing path, the receiving or destination router processes the packet to compare the address header information to the routing table maintained on the router for the next hop to forward the information packet. The router then forwards the information packet to the appropriate router on the next hop determined by the topological data in the routing table.

Packet Processing

Some routers are comprised of two types of components. The first type of component is the network processor. Network processors handle the basic routing functions by quickly processing (sometimes called “parsing”) address header information in the information packet and then forwarding the information packet to the next hop destination as specified by an address stored in a route cache entry in the processor. This lower-level processing and forwarding of an information packet is sometimes referred to as the “fast-path” processing.

The second type of component in a router is signaling processors.

If an information packets header address is not found in the route cache, higher-level processing may be required for matching the address with the routing table and forwarding the packet from the signaling processor. Additionally, some information packets require more extensive processing due to security, quality of service, or other control functions that require processing and/or implementation prior to forwarding of the information packet. This higher-level processing and forwarding of information packets is sometimes referred to as the "slow-path" processing.

Increasingly, information packet processing demands are expanding as networks begin to adopt the IP protocol. There is a growing need to enable access edge to access edge signaling for ubiquitous mobility, poke holes in firewalls, agree on a common method of dynamic quality of service (QoS) semantics (both per-flow and aggregate), accommodate security mechanisms, and support topological updates. These additional needs and accompanying processing demands on a network require increased scrutiny and processing of information packets using the slow-path processing technique. Consequently, the speed of information packet transmission and efficiency of communication decreases on the network because of the slow-path processing required for information packets on a

network. There is a need for a mechanism to segregate information packets requiring increased scrutiny in the signaling processor and information packets that can continue to be routed using the route cache of the network processor.

5

SUMMARY OF THE INVENTION

The present invention defines a filtered router alert hop-by-hop option associated with an information packet that designates the information packet as being of potential interest to the router requiring closer examination. The filtered router alert hop-by-hop option includes an identifier and a filtered router bitmap flag. The filtered router alert identifier informs the router that the information packet may require routing using the signaling processor.

The filtered router bitmap flag determines whether the transit router has a requirement for the data in the information packet and if the transit router must use the signaling processor. Information packets not having a filtered router alert hop-by-hop option identifier or a bitmap flag designating a requirement are processed and forwarded using the network processor.

20

BRIEF DESCRIPTION OF THE DRAWINGS

The objects and features of the invention will become more readily understood from the following detailed description and appended claims when read in conjunction with the accompanying drawings in which like numerals represent like elements and in which:

Fig. 1 is a schematic diagram of a packet-based communication system;

Fig. 2 is a general schematic diagram of an information packet used in packet-based communication systems;

Fig. 3 is a schematic diagram of a packet-based communication system showing routing for an information packet using the prior art;

Fig. 4 is the preferred format for a filtered router alert hop-by-hop option for use in IP version 6 (IPv6); and

Fig. 5 is a schematic diagram of a packet-based communication system showing routing for an information packet using the invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figure 1 is a schematic diagram of a packet-based communication system. Host 1 (H1) 5 is a computer, computer server, cellular phone, or

other communication device linked to communication network 80 by communication link 55. Host 2 (H2) 95 is a computer, computer server, cellular phone, or other communication device linked to communication network 80 by communication link 71. In more general terms, Host 1 5 and Host 2 95 are nodes in a network. Nodes are connection points, either a redistribution point or an end point for data transmissions having a destination address corresponding to that assigned to the node. Nodes possess the capability to process and forward information packets to other nodes. Hosts and routers, in general terms, are both considered to be nodes.

The communication network 80 is comprised of seven routers with a current topology of communication links 11, 13, 21, 22, 23, 31, 61, 51, 55, and 71. Router 1 (R1) 10 is connected to Router 2 (R2) 20 by communication link 13 and to Router 4 (R4) 40 by communication link 11. Router 4 (R4) 40 is connected to Router 2 (R2) 20 by communication link 41. Router 2 (R2) 20 is connected to Router 5 (R5) 50 by communication link 21, to Router 6 (R6) 60 by communication link 22, and to Router 3 (R3) 30 by communication link 23. Router 3 (R3) 30 is connected to Router 6 (R6) 60 by communication link 31. Router 6 (R6) 60 is connected to Router 5 (R5) 50 by communication link 61. Router 5 (R5) 50 is connected to Router 7 (R7) 70 by communication link 71. H1 5 is linked

to R1 10 by communication link 55, and H2 95 is linked to R7 70 by communication link 71. H1 5 and H2 95 can reside on the same network or different networks, and communication links 55 and 71 can include one or more intervening networks, including the Internet, or network 80 can be
5 part of the Internet.

Each of the transit routers R1 10, R2 20, R3 30, R4 40, R5 50, R6 60, and R7 70 can use either fast-path or slow-path routing techniques. Typically, information packets received at a transit router are processed only using the fast-path to retrieve address header data. However, some
10 information packets transiting network 80 require an increased level of examination to retrieve additional data and are processed using the slow-path technique. The use of the slow-path routing technique on every transit router significantly slows down transmission speed through network 80.

15 Figure 2 shows the general format of an information packet used in packet-based communication networks. Information packets use an encoding format of "1" and "0" data bits to build a data stream that a computer or other communication device can interpret. The information packet 100 has an address header (AH) 110 that provides routing instructions for transport over a packet-based communication system. The ad-
20

dress header includes data for the destination device and the originating device. The actual length and configuration of the address header 110 (AH) is dependent on the actual communication protocol being used in a given network's protocol implementation (e.g. IPv4, IPv6, etc).

5 The information packet 100 also contains a variable length data field (DF) 120 that contains the actual information being transmitted from the originating device to the destination device. Address data in the address header 110 can be retrieved by routers using fast-path processing, but other data within the information packet may require retrieval by
10 routers using the slow-path processing technique.

Figure 3 shows the routing path for an information packet transmitted over a communication network from Host 1 (H1) 205 to Host 2 (H2) 295. H1 205 is a computer, computer server, cellular phone, or other communication device linked to communication network 280 by communication link 255. H2 295 is a computer, computer server, cellular phone, or other communication device linked to communication network 280 by
15 communication link 271. Router 1 (R1) 210 is connected to Router 2 (R2) 220 by communication link 213 and to Router 4 (R4) 240 by communication link 211. Router 4 (R4) 240 is connected to Router 2 (R2) 220 by
20 communication link 241. Router 2 (R2) 220 is connected to Router 5 (R5)

250 by communication link 221, to Router 6 (R6) 260 by communication
link 222, and to Router 3 (R3) 230 by communication link 223. Router 3
(R3) 230 is connected to Router 6 (R6) 260 by communication link 231.
Router 6 (R6) 260 is connected to Router 5 (R5) 250 by communication
5 link 261. Router 5 (R5) 250 is connected to Router 7 (R7) 270 by com-
munication link 271. H1 205 is linked to R1 210 by communication link
255, and H2 295 is linked to R7 270 by communication link 271.

An information packet transmitted from H1 205 to H2 295 is
routed over the communication network 280 by a set of routers reflecting
10 the current topological configuration of the network 280 and the most effi-
cient and/or available path chosen for the packet by algorithms on each
router during the hop-by-hop transmission.

If all of the network routers use the slow-path routing technique,
the information packet is transmitted by link 255 from H1 205 to R1 210,
15 where the packet is processed via the slow-path processing technique.
From R1 210, the information packet is transmitted to R2 220 where it is
routed via the slow-path processing technique. From R2 220, the informa-
tion packet is forwarded to R6 260 where the packet is again processed via
the slow-path processing technique on R6 260. The packet is then trans-
20 mitted to R5 250 where the same slow-path processing technique is fol-

lowed. From R5 250, the information packet is transmitted to R7 270 where the information packet is again processed using the slow-path processing technique before being forwarded to the destination address of H2 295.

5 All five transit routers use the slow-path processing technique even though only two routers (e.g. R2 220 and R6 260) may actually require the additional data in the information packet extracted by the slow-path processing technique. Compared to the transmission time if all five transit routers used the fast-path processing technique, the transmission time for
10 the information packet over network 280 in Figure 3 is increased considerably as all five transit routers use the slow-path processing technique.

Figure 4 shows the format of an information packet for a filtered router alert hop-by-hop option of the invention for use with Internet Protocol version 6 (IPv6). A Protocol Data Unit (PDU) used in IPv6 incorporating the invention is defined. The PDU is comprised of 8-bytes (64-bits)
15 of data organized into four data fields. The Option Type (T) data field 305 designates the type of PDU option. The value of the first two bits is "0". This indicates to routers that do not recognize the option to forward the information packet to the next hop. The next bit value "1" indicates that
20 portions of the contents of this option are mutable. The value reflected by

the final five bits in the Option Type field 305 designates the PDU as a Filtered Router Alert Hop-by-Hop Option (the value 6 is a suggested designated value only).

5 The next data field is the Option Length (L) data field 310 which specifies the length of the PDU in bytes. The Option Value (V) data field 315 is a two-byte long data field. The Option Value data field 315 designates the type of communication protocol (e.g. message and Resource Reservation Protocol (RSVP) message) and/or the functional reason for the alert. This Option Value field 315 is processed on the fast-path to indicate to the router what specific parts of the information packet to examine more closely. For example, the V field 315 may direct increased examination of both mobile binding information and bandwidth reservation information which are found in different portions of the same information packet. Examples of value indicators include indicators for an information packet containing a Path Directed - Encapsulating Security Payload (PD-ESP) header, a mobility header, a Multicast Listener Discovery Protocol (MLDP) request, and/or a resource request. The Option Value field 315 can also be used to statistically multiplex the packet to an embedded function or to some external process specifically designed to handle the particular piece of functionality.

10

15

20

The Bitmap (B) data field 320 is four bytes long and is integral of the present invention. The Bitmap data field 320 contains a set of filtered router bitmap flags that indicate the type of relevant data within the information packet. There are 32 available filtered router bitmap flags in B
5 field 320. Nine bitmap flags (23-31) are defined for the B field 320, but additional flags are possible.

The "E" flag 321 switches the information packet to the slow-path for further inspection when the router being traversed is a Multi Protocol Label Switching (MPLS) label edge router. The "E" flag 321 is used to
10 trigger requests in the MPLS signaling domain.

The "G" flag 322 switches the information packet to the slow-path for further inspection when the router being traversed is acting as a security gateway. The "G" flag 322 is used with nodes of a Virtual Private Network (VPN) for signaling with the security gateway when the exact
15 address is not known or the discovery is deemed too inefficient.

The "N" flag 323 indicates slow-path processing for further inspection when the router being traversed is being used as a Network Address Translation (NAT) node. The "N" flag 323 is used by NAT administration tools as well as applications which require signaling for proper
20 NAT traversal.

The "C" flag bit 324 indicates directing to the slow-path for further inspection when congestion or load threshold is detected on an interface.

The "C" flag bit 324 is used by network analysis tools and traffic engineering applications.

5 The "A" flag bit 325 indicates directing to the slow-path when an aggregation function is provisioned for an interface. An example is a Wide Area Network (WAN) interface connecting multiple sites. The "A" flag 325 is used by network analysis tools and traffic engineering applications.

10 The "P" flag bit 326 switches the information packet to the slow-path for further inspection when a per-flow function is provisioned for an interface. Such an interface may exist between an over-provisioned Local Area Network (LAN) and WAN or any bandwidth-constrained shared link. The "P" flag bit 326 is used by network analysis tools and per-flow
15 resource management.

 The "S" flag bit 327 indicates slow-path routing on interfaces that entail a change in security keys. This bit flag may be used to indicate a difference in administrative domain ownership between peers or a downstream-upstream provider boundary. The links and interfaces between
20 boundary routers of separate Border Gateway Protocol (BGP) autonomous

systems are examples of where the "S" flag bit 327 is employed. The "S" flag bit 327 is used by network layer functions that re-use the existing security associations and trust relationships set up between systems.

5 The "M" flag bit 328 indicates slow-path routing is requested for an information packet on an interface which constitutes a layer 3 mobility-enabled edge router. An example for such a router is one close to the mobile device performing local mobility management functions or a router closer to the correspondent performing mobility functions.

10 The "F" flag 329 is employed to switch the information packet to slow-path routing when the router is an interface configured to access control functions such as in conjunction with a firewall.

Routers that recognize this Filtered Router Alert Hop-by-Hop Option during fast-path routing will recognize the applicable bit flag and transmit the packet to the signal processor for examination of the bit flag.

15 The network processor only needs to process the information packet in sufficient detail to determine whether the information packet contains data of interest to the router requiring more detailed examination and slow-path processing. The filter flags in the B field 320 provide a quick method during fast-path routing to determine the appropriate level of interest by the

20 network processor based upon an examination of the filter flags. Once the

network processor determines slow-path routing is requested, the information packet is forwarded onto the slow-path for more detailed examination and processing.

If the field values become corrupted during transmission, the router
5 can mute the filter flag field. That is, the filter flags are mutable in that a network processor can ignore the filter flags after an integrity check shows some data corruption on the information packet. This mutability enhances the flexibility and scalability of applications using the Filtered Router Alert Hop-by-Hop Option.

10 Applications that can benefit from the Filtered Router Alert Hop-by-Hop Option include congestion avoidance mechanisms, communications with NAT and firewall devices, per-flow resource management, aggregate resource management, and network security association establishment.

15 Congestion avoidance mechanisms can benefit from the knowledge of congestion at specific interfaces. The "C" flag 324 can be used to detect, act on, and report on this congestion. Such tools can be used by network administrators to determine the points of congestion within networks on a real-time basis.

Some applications can benefit from knowledge of and communication with NAT or firewall devices. The “N” flag 323 and the “F” flag 329 can be used to facilitate signaling to these applications. Tools can also perform network analysis related to address translation technology.

5 Networks and applications can also benefit from a common method to signal per-flow resource management functions. The “P” flag 326 facilitates signaling per-flow resource management functions on interfaces performing aggregate resource management.

10 Network security imposes a growing requirement for network nodes to quickly exchange and cache shared secret keys and establish security associations with other nodes along a bearer path at varying levels of granularity. Current mechanisms for performing these functions force an unpredictable delay by having to either resort to using trusted authentication service or by exchanging keys dynamically on a per-use basis. This
15 can cause an information packet to be queued while waiting for authorization and result in the stateful queuing of an information packet causing discontinuity in the flow (e.g. queue and wait) with arbitrary delay.

 The “S” flag 327 allows key exchange and security associations to be established by leveraging the trust relationship and security associa-
20 tions of the routing system itself, facilitating the establishment of security

associations between an arbitrary set of endpoints along a path using a single roundtrip message exchange. The PD-ESP value in conjunction with the "S" flag 327 can be used to establish this secure routing path for a communication.

5 Figure 5 shows the routing path for an information packet transmitted over a communication network using the present invention. H1 405 is a computer, computer server, cellular phone, or other communication device linked to communication network 480 by communication link 455. H2 495 is a computer, computer server, cellular phone, or other communication device linked to communication network 480 by communication
10 link 471. R1 410 is connected to R2 420 by communication link 413 and to R4 440 by communication link 411. R4 440 is connected to R2 420 by communication link 441. R2 420 is connected to R5 450 by communication link 421, to R6 460 by communication link 422, and to R3 430 by
15 communication link 423. R3 430 is connected to R6 460 by communication link 431. R6 460 is connected to R5 450 by communication link 461. R5 450 is connected to R7 470 by communication link 471. H1 405 is linked to R1 410 by communication link 455, and H2 495 is linked to R7 470 by communication link 471. H1 405 and H2 495 can reside on the

same network or different networks, and communication links 455 and 471 can include one or more intervening networks, including the Internet.

An information packet transmitted from H1 405 to H2 495 is routed over the communication network 480 by a set of transit routers reflecting the current topological configuration of the network 480 and the most efficient and/or available path chosen for the packet by algorithms on each router during the hop-by-hop transmission. Because of the network 480 service requirements, information packets are used to communicate data that is only recovered during slow-path routing of the information packet. In this example, data in the information packet found during slow-path routing is only needed in R2 420 and R6 460.

Using an information packet containing a Filtered Router Alert Hop-by-Hop Option, the information packet is transmitted from H1 405 to R1 410, where the packet is processed using the fast-path processing technique. R1 410 does not need slow-path processing, so none of the bitmap flags in the information packet match any of the flags in the provisioned data field on R1 410. Because slow-path processing is not requested, the fast-path processing technique is used in R1 410. The information packet is forwarded to R2 420 only using fast-path routing.

R2 420 requires slow-path processing for certain information packets. As the information packet is processed on router R2 420, the Filtered Router Alert Hop-by-Hop Option is encountered. The filtered router bitmap flags is compared to the provisioned data field on the router R2
5 420, which identifies the information packet as one that needs increased examination under slow-path processing by the signal processor. The information packet is then forwarded to the signal processor for slow-path processing.

After being examined and processed as required, the information
10 packet is forwarded to R6 460. R6 460 also requires slow-path processing for certain information packets. The packet is examined and the Filtered Router Alert Hop-by-Hop Option is encountered, identifying the information packet as requiring increased examination and processing by the signal processor. The information packet is forwarded to the signal processor
15 on R6 460 for use as defined by the bit flag values. The information packet is then transmitted from R6 460 to R5 450 for further routing. Because R5 450 has no requirement for slow-path processing, none of the filtered router bitmap flags in the information packet match any of the flags in the provisioned data field on R5 450. The information packet is

forwarded to R7 470, where it is again processed using the fast-path processing technique before forwarding to the destination address of H2 495.

As specified in the present invention, slow-path routing is selectively conducted on only routers R2 420 and R6 460 of the network. Instead of requiring slow-path routing on each of the five transit routers,
5 only two of the five transit routers use slow-path routing techniques using the invention - and the other three transit routers continue to use fast-path routing techniques. By reducing the number of transit routers required to use the slow-path processing technique, the transit time over the network
10 480 is decreased considerably compared to the transit time found using all slow-path routing techniques.

While the invention has been particularly shown and described with respect to preferred embodiments, it will be readily understood that minor changes in the details of the invention may be made without departing from the spirit of the invention. Having described the invention, I
15 claim: